arXiv:1608.03024v1 [stat.AP] 10 Aug 2016

A Spatial Generalized Linear Mixed Model for Analyzing Efficiency of Foreign Aid Allocation

Philip A. White*[1], Candace Berrett†[2], E. Shannon Neeley-Tass‡[2], and Michael G. Findley§[3]

[1]*Department of Statistics, Duke University, Durham, NC, USA*

[2]*Department of Statistics, Brigham Young University, Provo, UT, USA*

[3]*Department of Government, University of Texas at Austin, Austin, TX, USA*

August 11, 2016

Abstract

The Open Aid Malawi initiative has collected an unprecedented database that identifies as much location-specific information as possible for each of 2500 individual foreign aid donations to Malawi since 2003. Because so much foreign aid is being sent to Malawi, ensuring efficient use and distribution of that aid is important to donors and to Malawi citizens. However, because of individual donor goals and difficulty in tracking donor coordination, determining presence or absence of efficient aid allocation is difficult. We compare various Bayesian spatial linear regression models to relate aid allocation to various economic indicators within seven donation sectors. We find that the spatial gamma regression model best predicts current aid allocation. From this model, we identify evidence both in favor and against efficient aid allocation. Thus, we develop a model to create a more efficient aid allocation scenario; comparing this to the current allocation provides insight for improving future aid allocation.

Keywords: gamma regression, Bayesian, posterior predictive, donor collaboration, Malawi

*philip.a.white@duke.edu
†cberrett@stat.byu.edu
‡tass@stat.byu.edu
§mikefindley@utexas.edu

1 Introduction

Each year, wealthy countries and international organizations send billions of dollars in foreign aid to developing countries across the world. The Organization for Economic Co-operation and Development (OECD) estimates that in 2015 the United States alone donated 31.1 *billion* dollars in development assistance and that other developed countries are donating anywhere between 700 million and 19 billion US dollars (OECD, 2015). Aid giving is thus big business, and we need to understand better whether the business operates in the service of its so-called clients in recipient countries or rather in service of its own interests.

Donors' motives are varied, ranging from pursuing strategic interest (Alesina and Dollar, 2000) to trying to facilitate positive social and economic outcomes. And yet, with so much spent on development assistance, donors, governments, and citizens may want these donations to be allocated efficiently — providing money to places and situations with the most need. Even where strategic interest concerns are high, and therefore aid may not be allocated with need in mind, foreign aid can affect peoples' lives in non-trivial ways, for better or for worse. Thus, stakeholders in both recipient and donor countries deserve to know whether aid is in fact allocated in ways that respond to and meet the needs of recipient countries. For example, donations meant for improving education should ideally be sent to those places within the country where education is suffering the most.

Several initiatives have been introduced to improve aid targeting and coordination, including the Paris Declaration of 2005 which called for better alignment between government objectives and donor aid allocation (Aced, 2009; Findley et al., 2016). Communication among donors and recipients, as well as donors and other donors, in order to track and efficiently allocate the vast array of projects has proven nearly impossible, however. Moreover, there has been a dearth of scholarly attempts to track the complexity of aid allocation and so both academics and policy makers have a poor understanding of the patterns of foreign aid allocation. Because so much hangs in the balance, answers to the question of efficient aid allocation are sorely needed.

In 2011, the government of Malawi, together with the University of Texas at Austin's Cli-

mate Change and African Political Stability team, began efforts to collect and display donor information at the subnational level (Findley et al., 2011; Weaver et al., 2014). This initiative — Open Aid Malawi — resulted in an unprecedented and unique database containing as much location-specific information as possible for *each* individual donation to Malawi since 2003 (Peratsakis et al., 2012). This database contains over 2500 tracked donations all at the subnational level, to date the largest of its kind. Such a large and detailed resource makes it possible to examine donor aid across various project sectors.

Malawi is a particularly interesting country for examining aid allocation. It is located in Southern Africa and is populated by over 16 million people. Malawi is divided into 3 regions and 28 districts. Gaining its independence in 1964, the country is one of the youngest, most diverse, and densely populated places in the world. Malawi's economy, education, and public health, however, are some of the least developed anywhere in the world. Specifically, they have an infant mortality rate of 69.3 for every 1,000 births, 50.7% of the population live below the poverty line, a 68.2% literacy rate, and agricultural goods are their leading export (National Statistical Office, 2015). In an effort to improve circumstances in Malawi, a large number of foreign donors support economic, health, and educational development and improvement plans in Malawi. In contrast to most highly impoverished countries, Malawi has experienced almost no violence, making it an ideal location to study aid allocation patterns as the presence of violence makes stable aid allocation patterns difficult. Put differently, given its political stability Malawi is a most likely case to observe donors responding to need, and we thus explore in this paper the extent to which donors in fact allocate based on indicators of need.

Because so much aid is sent to Malawi and because of its need for hastened development, the question of whether aid is allocated efficiently is paramount both for taxpayers in donor countries as well as citizens affected by aid in recipient countries. We are specifically interested in quantifying how assistance is being allocated within the country to determine whether sectoral aid matches sectoral need. In particular, we focus on seven project sectors or donation types: (1) Agriculture; (2) Education; (3) Governance; (4) Health; (5) Rural Development [RD]; (6) Roads, Public Works, and Transportation [RPT]; and (7) Water, Sanitation, and Irrigation [WSI].

Using district-specific economic, social, and demographic data (hereafter referred to as economic indicators for short), we model foreign aid allocation in all districts for these project sectors. While we do not expect the economic indicators to be the sole driving forces behind aid allocation, by donors own admission they should be key factors in determining whether or not aid is being allocated efficiently. Therefore, our goals are, first, identifying high-level economic indicators that are related to aid allocation while accounting for and making use of spatial dependence; second, determining the ability of our model to capture a true relationship between aid distribution and the economic indicators; and third, to propose an efficient allocation scenario and its discrepancy from the actual allocation in order to inform future donations.

Analyses attempting to measure donor coordination in Malawi using this novel data set have begun. For example, comparing donations before and after the Paris Declaration of 2005 with various economic and demographic indicators, Nunnenkamp et al. (2015) found no evidence of increased donor coordination across sectors. De and Becker (2015) used instrumental variable techniques to show evidence of improved economic conditions at the local level from aid donations, emphasizing the need and ability for donors to coordinate their efforts using local economic indicators. Using a model to accurately capture the skewness present in aid data, we seek to significantly expand on these analyses by exploring the extent of the spatial dependence of the different sectors and to account for this spatial dependence to make inferences about relationships between donations and economic indicators at the local level.

In this manuscript, we examine donation allocation within Malawi at the subnational level using a Bayesian spatial gamma regression model. We describe the data in more detail in Section 2. In Section 3 we explicitly define the model used for analysis as well as the comparison models. In Section 4.1 we provide an in-depth analysis of the various model parameters and their interpretations and implications for aid allocation in Malawi, contrasting these results with a more ideal allocation scenario in Section 4.2. Finally, we provide concluding remarks in Section 5.

2 Data

The Open Aid Malawi database (Peratsakis et al., 2012) contains over 2500 donations spanning the years 2003 to 2016. In addition to the amount of each donation in US dollars (USD), the database also includes information about each donation's intended project sector (Agriculture; Education; Governance; Health; Rural Development; Roads, Public Works, and Transport; and Water, Sanitation, and Irrigation), spatial location, and "precision" or level at which a donation was designated (e.g., country, region, district, city). Figures 1a-1d illustrate the different precisions at which a donation could be assigned. For example, Figure 1b shows the total money per person assigned to the three regions (North, Central, South) in Malawi, while Figure 1d shows the locations of the donations assigned to a city (blue star points) or a more specific latitude-longitude (red circle points).

In addition to the donation data, we also have population and economic indicator data for each district from the 2008 Malawi census (National Statistical Office, 2008). These economic indicators include: proportion of people in poverty, average distance of people to schools (km), number of recently (within two weeks) injured people, food spending as a percentage of all expenses, mean area cultivated per household, and proportion of population with electric lighting. Descriptions of the variables are in Table 1. However, we do not have these values for the three smallest districts – Likoma (a small island in Lake Malawi), Neno (split in 2003 from Mwanza), and Area under National Administration (various locations across the country). While initially we imputed these missing values using the expectation-maximization (EM) algorithm, we ultimately decided not to include these districts in our analysis. This did not have noticeable impact on the results.

Because our economic indicator data is only available at the district level, we aggregated all the donation data to the district level. For the country- and region-level donations, we assumed these were dispersed proportionally among the districts according to district population. Our final donation data is therefore the total donation sum for each district in USD by project sector. Figure 1e shows the total donations summed across project sectors. We expect that more money should be sent to those districts with more people; thus, we use dollars per person

as the response variable by dividing the district donation total for each project sector by the district population. For the final derived dataset, the total sample size is $N = 182$ since the data contain dollars donated per person in 26 districts for seven project sectors. Note that, as shown in Figure 1f, the data is right-skewed, strictly positive, and can be mapped over Malawi.

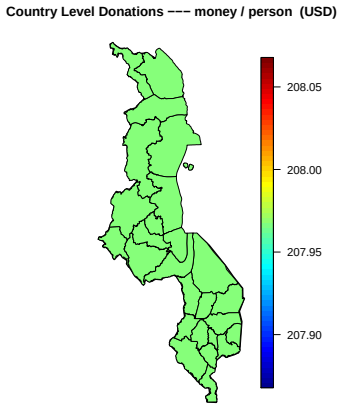

(a) Country Level Donations

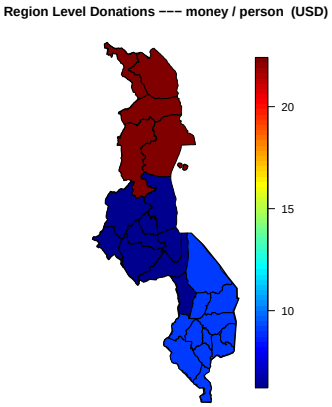

(b) Region Level Donations

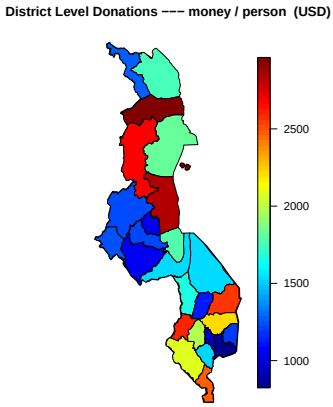

(c) District Level Donations

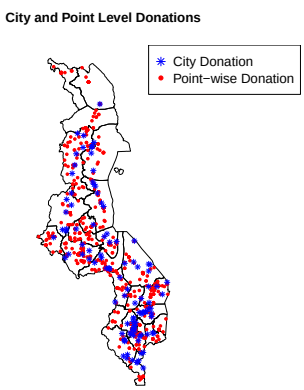

(d) Point Level Donations

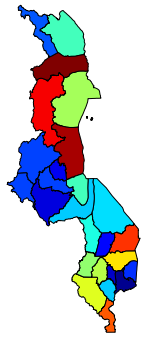

(e) Summed Donations

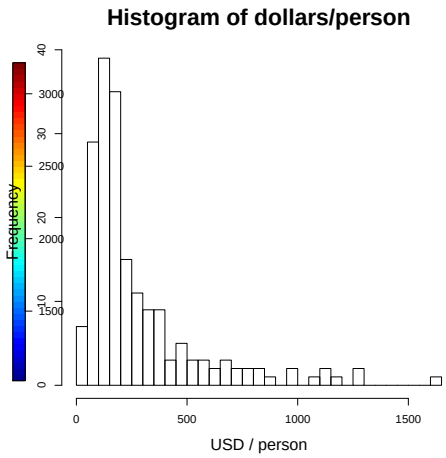

(f) Histogram of Donations

Figure 1: Money per person for donations at the various levels of location information available. Figures (a) – (c) show maps of money per person for (a) donations made to the entire country, (b) donations made to the individual northern, central, and southern regions, and (c) donations made to the individual districts. Figure (d) shows a map of the locations of city and point-wise donations. Figure (e) shows the total money per person within each district (total money per person from (a) – (d) within each district). Finally, (f) shows a histogram of the donations per person within each project sector and district.

Variable	Abbreviation	Description
Sectors		
Agriculture	Agriculture	Donations assigned to development of agriculture
Education	Education	Donations assigned to education development
Governance	Governance	Donations assigned to both economic and democratic governance
Health	Health	Donations assigned to for development of health initiatives
Integrated Rural Development	RD	Donations assigned for rural development
Roads, Public Works, and Transportation	RPT	Donations assigned to developing roads, transportation, and other public works
Water, Sanitation, and Irrigation	WSI	Donations assigned to water, sanitation, and irrigation development
Economic Indicators		
Proportion in Poverty	Poverty	Percent of population living below the national poverty line
Distance to School	Dist to School	Average distance (km) across households to the nearest school in each district
All Injured	All Injured	Number injured or ill in the previous 2 weeks (as of the 2008 census) in each district
Food as % of Expenditures	Food %	Average percent of household expenditures spent on food in each district
Mean Land Cultivated per House	Land Cultivated	Average area cultivated per household (in hectres) in each district
Electric Lighting Proportion	Electric %	Percent of households in the district with electric lighting
Population	Population	Population of each district

Table 1: Names and descriptions of the data variables used in the analysis.

3 Model

We expect some quantities, like agricultural output, public health outcomes, or non-urban areas, to be grouped spatially; therefore, it is plausible that accounting for spatial dependence could improve model performance in terms of prediction and inferences. The normal linear regression model with spatial random effects (a Gaussian process model) is a standard model for spatially-distributed quantities (Cornford et al., 2005; Banerjee et al., 2008). However, because our data is strictly positive and right-skewed, we are interested in more general models. Following Diggle et al. (1998), we make use of a spatial generalized linear mixed model (SGLMM).

Let Y_{ij} be the dollars per person (USD) for the i^{th} district ($i = 1, \ldots, 26$) and the j^{th} project sector ($j = 1, \ldots, 7$). We model

$$Y_{ij}|\mu_{ij}, \theta \sim f(y_{ij}|\mu_{ij}, \theta),$$

where $\mu_{ij} = E(Y_{ij}|\mu_{ij}, \theta)$ and θ is an optional scale parameter. We model the mean using a link function, $g(\cdot)$,

$$g(\mu_{ij}) = \mathbf{x}'_{ij}\boldsymbol{\beta}_j + \phi_{ij}, \tag{1}$$

where $\mathbf{x}_{ij}$ is a $k \times 1$ vector of covariates for the i^{th} district and the j^{th} project sector, $\boldsymbol{\beta}_j$ is a $k \times 1$ vector of coefficients for each sector, and ϕ_{ij} is a spatial random effect where each sector has its own spatial dependence structure. We model the spatial random effect using

$$\boldsymbol{\phi}_j \sim \mathrm{N}(\mathbf{0}, \sigma^{\mathbf{2}}\mathbf{R_j}),$$

where $\boldsymbol{\phi}_j$ is the 26×1 vector of random effects for the j^{th} project sector, σ^2 is a variance parameter constant across sectors, and $\mathbf{R}_j$ is a 26×26 spatial correlation matrix unique for

each sector. The lm^{th} element of $\mathbf{R}_j$ is defined by the Matérn correlation function,

$$\frac{1}{\Gamma(\nu)2^{\nu-1}}\left(\frac{d_{lm}}{\rho_j}\right)^{\nu} K_{\nu}\left(\frac{d_{lm}}{\rho_j}\right),$$

where d_{lm} is the distance (in km) between the centers of districts l and m. We choose the Matérn correlation function because it provides a general covariance function that can approximate diverse semivariograms by adjusting its (site-specific) range parameter, ρ_j, and smoothness parameter, ν (Handcock and Stein, 1993; Minasny and McBratney, 2005).

Because we expect a combination of the economic indicators to be related to aid allocation within each sector, we fit a sector-specific coefficient for each of the seven economic indicators (including population) as well as a sector-specific intercept to account for the different amounts of aid appointed to each sector. Therefore, there are $k = 8$ coefficients for each sector, or 56 total coefficients in our model.

Finally, we use relatively non-informative prior distributions. Specifically,

$$\begin{aligned}
\theta &\sim \text{Inverse-Gamma}(a_\theta, b_\theta) \\
\boldsymbol{\beta} &\sim \text{Normal}(\mathbf{0}, s^2\mathbf{I}) \\
\sigma^2 &\sim \text{Inverse-Gamma}(a_\sigma, b_\sigma) \\
\rho_j &\sim \text{Gamma}(a_\rho, b_\rho) \\
\nu &= 0.5
\end{aligned}$$

where $a_\theta = 1$, $b_\theta = 1$, $s^2 = 10^9$, $a_\sigma = 1$, $b_\sigma = 1$, $a_\rho = 5$, and $b_\rho = 0.05$. We select these values for the hyperparameters so that our prior distributions are diffuse and so that the *a priori* mean effective range ρ_j is 100 km. In a simulation study, we found that the model was not sensitive to the prior parameter values.

3.1 Model Selection

We consider four Bayesian models, comparing both spatial and non-spatial models using gamma, negative binomial, and log-normal likelihood distributions (see Table 2). For the

Model	Likelihood: $f(y_{ij}\|\mu_{ij},\theta)$	Link: $g(\mu_{ij})$	Spatial Dependence
1	Gamma	$\log(\mu_{ij})$	No
2	Gamma	$\log(\mu_{ij})$	Yes
3	Negative Binomial	$\log(\mu_{ij})$	No
4	Negative Binomial	$\log(\mu_{ij})$	Yes
5	Log-normal	μ_{ij}	No
6	Log-normal	μ_{ij}	Yes

Table 2: Models used for comparison.

non-spatial models, we assume $\phi_{ij} \equiv 0$ for all i and j in equation (1). We fit each model using a Gibbs sampler with 100,000 post-burn-in iterations (and 20,000 draws discarded as burn-in). Because of the model formulation, we are able to use closed form complete conditional distributions for $\boldsymbol{\beta}$ and σ^2, while we used the Metropolis-Hastings algorithm for all other model parameters.

We are interested in the model that performs the best in predicting foreign aid for each sector in each district; therefore, we use the posterior predictive distribution,

$$f(\boldsymbol{y}_{\text{new}} \mid \boldsymbol{y}) = \int_{\Psi} f(\boldsymbol{y}_{\text{new}} \mid \boldsymbol{\psi})\, \pi(\boldsymbol{\psi}|\boldsymbol{y})\, \mathrm{d}\boldsymbol{\psi}, \tag{2}$$

where $\boldsymbol{\psi}$ is the set of all model parameters and Ψ is its parameter space, for each model to determine which model most accurately predicts foreign aid in Malawi. We consider both mean squared error (MSE) and mean absolute error (MAE) between the posterior predictive mean and observed donations. We use these measures since there are some anomalously high values of y_{ij} and we are interested in selecting the model that predicts the majority of these values accurately.

Table 3 shows the MAE and MSE for the in-sample predictions from each of the models. The spatial models drastically outperform the non-spatial models, providing evidence that the spatial dependence is providing additional information about the allocation of foreign aid beyond the information provided by the economic indicators. Model 2, or the spatially-dependent gamma model, provides the smallest MAE and MSE. The next best model is Model 6, or the spatially-dependent log normal model; it's MAE and MSE are relatively small as well, but are more than 18 times as large as the errors from the spatial gamma model. While log nor-

	Model	MAE (USD)	MSE (USD^2)
1	Non-Spatial Gamma (log link)	187.4	109927.5
2	**Spatial Gamma (log link)**	0.030	0.007
3	Non-spatial Negative Binomial	241.0	122160.3
4	Spatial Negative Binomial	17.4	6678.6
3	Non-spatial Log Normal (identity link)	111.3	41033.5
4	Spatial Log Normal (identity link)	0.565	0.916

Table 3: Observed mean absolute error (MAE) and mean squared error (MSE) for the six posed model. Model 2, the spatially-dependent gamma model, is bolded because it performs best for both errors.

mal regression models tend to be preferred to gamma regression models due to simplicity in model-fitting (see, for example McCullagh and Nelder, 1989, pg. 286), in this case, because we are interested in the model that best reflects the data, we selected the spatial gamma model to use in our analysis.

4 Analysis

In this section, we examine the parameters and fitted values of the spatial gamma model described in Section 3. In Section 4.1, we discuss evidences for and against donor collaboration or efficient aid allocation by examining the posterior distributions of the model coefficients. We conduct an in-depth study in Section 4.2 to identify potential areas for improved aid allocation.

4.1 Parameter Estimates and Interpretation

Table 4 shows posterior means and standard deviations for all regression coefficients $\boldsymbol{\beta}$ as well as the marginal posterior probabilities that each regression coefficient is greater than 0. We use these posterior probabilities to examine the relationship between the economic indicator and the donations within a sector. Bolded values correspond to coefficients with posterior probability greater than 0.9 of being either positive or negative. The simulation study in Section 4.2 suggests these posterior probability cut-off values are reasonable (and even a bit stringent) for identifying relationships between economic indicators and aid per person.

Sector	*Agriculture*	*Education*	*Governance*	*Health*	*RD*	*RPT*	*WSI*
Covariate	mean (sd) $P(>0)$	mean (sd) $P(>0)$	mean (sd) $P(>0)$	mean (sd) $P(>0)$	mean (sd) $P(>0)$	mean (sd) $P(>0)$	mean (sd) $P(>0)$
Intercept	5.99 (1.22) 1.00	5.93 (1.30) 1.00	4.63 (1.11) 1.00	5.98 (1.12) 1.00	2.67 (1.66) 0.95	6.07 (1.37) 1.00	6.92 (1.46) 1.00
Poverty	0.01 (1.29) 0.51	**1.80** (1.37) **0.91**	0.12 (1.13) 0.54	0.31 (1.17) 0.61	**-2.17** (1.66) **0.09**	-0.99 (1.43) 0.24	-0.64 (1.54) 0.34
Dist to School	-0.01 (0.29) 0.49	**-0.46** (0.30) **0.06**	0.01 (0.25) 0.51	-0.12 (0.25) 0.32	**0.91** (0.39) **0.99**	0.28 (0.32) 0.81	**-0.43** (0.35) **0.11**
All Injured	-0.01 (0.02) 0.23	0.02 (0.02) 0.87	0.00 (0.02) 0.50	0.00 (0.02) 0.44	**-0.03** (0.02) **0.07**	-0.01 (0.02) 0.34	0.01 (0.02) 0.65
Food %	0.02 (0.68) 0.49	-0.75 (0.72) 0.14	-0.16 (0.60) 0.38	0.42 (0.61) 0.76	**1.92** (0.99) **0.97**	-0.10 (0.79) 0.45	-0.74 (0.84) 0.18
Land Cultivated	-0.24 (0.49) 0.31	-0.14 (0.51) 0.39	-0.06 (0.44) 0.44	-0.45 (0.45) 0.15	**0.83** (0.65) **0.91**	-0.42 (0.56) 0.22	0.40 (0.55) 0.77
Electric %	0.51 (0.54) 0.83	0.44 (0.57) 0.78	0.25 (0.47) 0.71	0.43 (0.49) 0.82	0.33 (0.71) 0.68	0.37 (0.61) 0.73	-0.30 (0.64) 0.31
Population ($\times 10^{-6}$)	**-0.63** (0.34) **0.03**	-0.29 (0.35) 0.20	0.01 (0.30) 0.52	**-0.51** (0.30) **0.04**	**0.80** (0.44) **0.96**	0.31 (0.38) 0.80	0.05 (0.40) 0.56

Table 4: Summary of the posterior distributions of the coefficients. '$P(>0)$' means the posterior probability that the corresponding coefficient is greater than 0. Coefficients (other than intercepts) that have >90% posterior probability of being either positive or negative are bolded. RD=Integrated Rural Development; RPT=Roads, Public Works, & Transportation; WSI=Water, Sanitation, Irrigation.

In large part, there is not evidence that the economic indicators are related to districts' aid per person within a sector at all. Many of the coefficient posterior distributions envelop 0 (those non-bolded coefficients). In fact, two sectors – Governance and Roads, Public Works, and Transportation (RPT) – appear to not be related to any of the economic indicators. For Governance, this isn't surprising as more than 50% of the donations from that sector were given at the country or region level (see Figure 1a and 1b), and we assigned these projects an equal weight per person to each of the districts within the country or region, respectively. For RPT, however, this isn't the case and a more likely explanation is just a lack of a strong relationship between aid and the economic indicators.

There are, however, several coefficients that provide evidence of efficient aid allocation. For example, distance to school has a 0.99 posterior probability of being positively related to RD aid. This makes sense since districts where the people are farther from a school tend to be more rural and thus would benefit from RD aid. Similarly, there is evidence that aid from the RD sector is also positively related to food expenditure and aid from the education sector is positively related to poverty, matching intuition for efficient allocation.

In contrast, distance to school is also related to Education aid, but its relationship is the opposite of what we would expect (only a 0.06 posterior probability of a positive relationship). In this case, the farther a districts' people are from a school, the less money for education that district receives. In fact, this appears to be the case for many of the covariates we would hope to be strongly related under an efficient aid allocation scenario. If there is not a strong contradictory relationship, there is not evidence of a relationship at all.

Obviously, these interpretations of coefficients are conditional on the values of the other covariates and multicollinearity among the covariates could introduce changes in the values of the coefficients when some covariates are left out of the model (thus changing interpretations). However, there is not much collinearity between the economic indicators, the highest absolute value correlation is 0.3, with most absolute value correlations less than 0.15. That said, interpretations of the coefficients must be made with care.

The spatial correlation in this model is governed by the district-specific range parameters ρ_j, which quantifies how donations vary spatially. The posterior distribution results are

displayed in Table 5. Importantly, the spatial correlation ranges are quite different between project sectors. RD and WSI have relatively short correlation ranges, meaning that these donation patterns change more quickly over space (see Figures 2c and 4e). In contrast, donations for Governance, Agriculture, and Health have relatively long correlation ranges, signaling that these donations are distributed more evenly or smoothly across Malawi (see Figures 2b, 4a, and 4c). Recall that many Governance donations were distributed equally across the country and region, and therefore the smooth distribution clearly makes sense for that sector. The other sectors' estimated effective ranges match intuition as well. We expect the need for aid allocation of RD and WSI to vary across Malawi, while we expect the need for aid allocation of Agriculture and Health to be more uniform over Malawi. Thus, the spatial correlation patterns for aid allow flexibility in how donations for different sectors vary over space.

	Posterior Mean	95% CI
Agriculture	112.00	(46.38, 202.49)
Education	93.33	(35.00, 170.86)
Governance	166.40	(69.56, 297.52)
Health	143.68	(59.60, 253.73)
IRD	31.65	(11.08, 62.95)
RPT	74.64	(31.10, 136.50)
WSI	65.25	(24.61, 125.70)

Table 5: Posterior summary for the spatial range parameters ρ_j for each sector j (reported in km).

Figure 2 shows plots of observed, posterior predictive means, and the residuals of total donations, and donations for Governance and RD (the plots for the remaining sectors are in the Appendix in Figure 4). These plots show how well our model is able to capture the district-level total and sector-specific donations, with a maximum residual of only 1.04 dollars per person. Much of this strength in prediction comes from the spatial random effect in the model, although the residuals still exhibit some spatial dependence even though they are small.

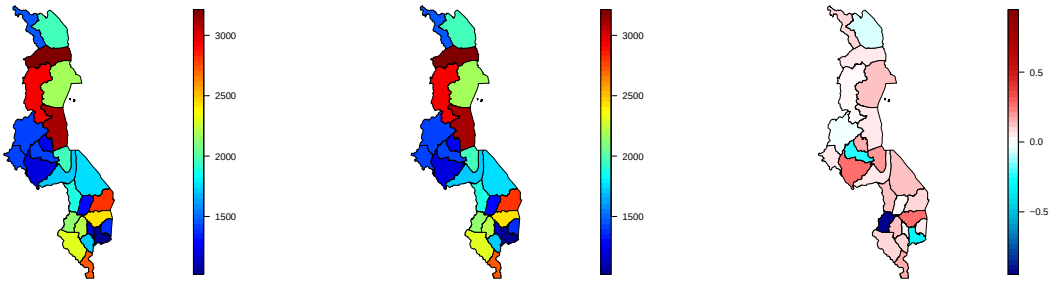

(a) Total Donation

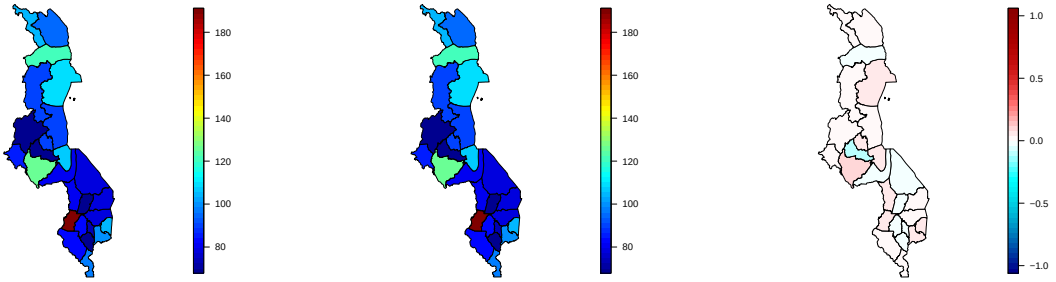

(b) Governance

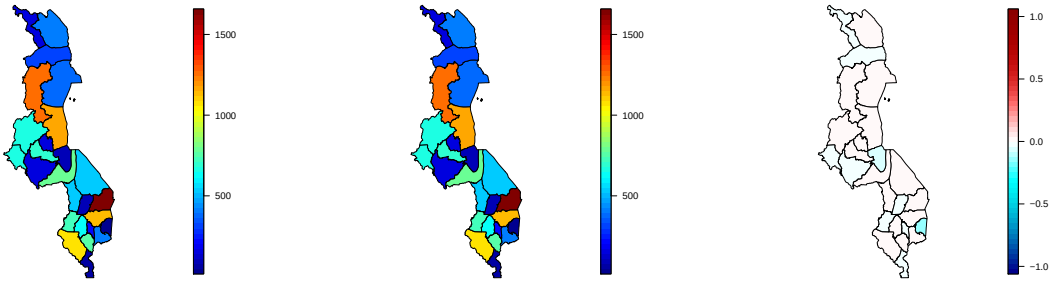

(c) Rural Development

Figure 2: Observed (left), predicted (center), and residual (right) donations plots for total donations (a) and the Governance (b) and RD (c) sectors. All values are in USD.

4.2 More Efficient Aid Allocation

Although we do not expect the economic indicators to be perfect drivers of aid allocation, they can provide direction for better aid allocation. In this section, we consider the scenario where aid is allocated according to a more expected or ideal relationship with specific economic

Sector	Economic Indicator	Expected Relationship
Agriculture	Land Cultivated	Positive
Education	Dist to School	Positive
Governance	Poverty	Positive
Health	All Injured	Positive
RD	Electric %	Negative
RPT	Electric %	Negative
WSI	Land Cultivated	Positive

Table 6: Selected economic indicator for each sector and the corresponding direction of the relationship.

indicators. We call this scenario our "efficient scenario." By doing this, we can first examine the potential for our model to capture a "true" relationship between aid and economic indicators. Second, and more important for donors as they identify locations to send aid, is to identify locations where more or less aid is needed within each sector.

To accomplish this task, we select a single economic indicator for each sector and identify its expected relationship to aid, shown in Table 6. We then use the posterior distributions from the fitted model of Section 4.1, but adjust the coefficients for these selected economic indicators. Specifically, we set the mean of the corresponding coefficient of each economic indicator of interest to be three posterior standard deviations above (or below for a negative relationship) zero. For example, the adjusted mean of the coefficient corresponding to Land Cultivated for Agriculture is set to 1.47 (see Table 4), while the other coefficients for the Agriculture sector are left the same. Then, the adjusted posterior distribution on the coefficients is

$$\boldsymbol{\beta}^* \sim N(\bar{\boldsymbol{\beta}}^*, \boldsymbol{\Sigma}(\boldsymbol{\beta})), \tag{3}$$

where $\bar{\boldsymbol{\beta}}^*$ is the adjusted posterior mean of $\boldsymbol{\beta}$ as described above and $\boldsymbol{\Sigma}(\boldsymbol{\beta})$ is the posterior variance of $\boldsymbol{\beta}$ from the fitted spatial gamma model in Section 4.1. The posterior correlation among the $\boldsymbol{\beta}$'s was quite small, and thus $\boldsymbol{\Sigma}(\boldsymbol{\beta})$ could also be a diagonal matrix with the marginal variances of each coefficient on the diagonal.

We then obtain adjusted posterior predictive distributions on aid per person for each sector using equation (2), but with the adjusted posterior distribution of $\boldsymbol{\beta}$. The computational

algorithm is as follows:

1. Draw a value of $\boldsymbol{\beta}^*$ from equation (3).
2. Use a draw of $\boldsymbol{\phi}$ and θ from the posterior draws obtained in Section 4.1.
3. Draw a value of the adjusted aid allocation, y^*_{ij}, for each i and j from the gamma likelihood, conditional on the values in steps 1 and 2.
4. Scale y^*_{ij} so that the new sector totals are the same as the data sector totals: $\sum_{i=1}^{26} y^*_{ij} = \sum_{i=1}^{26} y_{ij}$. (This requires that the total money within a sector does not change, only its distribution across the districts.)
5. Repeat for all post-burn-in posterior draws.

We can then use these draws from the adjusted posterior predictive distributions to make inferences on aid allocation.

To learn about the ability of our model to identify a "true" relationship, we refit the spatial gamma model to 50 randomly selected draws from the adjusted posterior predictive distribution. In this way we are making use of frequentist sampling distribution ideas to compute a power-type diagnostic. For these 50 randomly selected draws, using a 90% posterior probability of being greater than zero (or less than zero for expected negative relationships) cut-off, we would have identified a relationship 77% of the time. This is a bit lower than the ideal 90%, but this is likely due to re-scaling the data to keep total donation amounts the same. Notably, however, we never found a relationship in the opposite direction; i.e., a parameter that was truly positive never had a 90% posterior probability of being negative. Recall that Education was negatively related to distance to school – the opposite of what we would hope to see. This simulation study provides more evidence that if Education aid were actually positively related to distance to school, we would not have found this strong negative relationship, reinforcing the need for adjustments in aid allocation.

Finally, we compare the adjusted posterior predictive distributions of donation amounts to the observed donation amounts. Figure 3 shows standardized residual maps; specifically, for district i and sector j,

$$r_{ij} = \frac{y_{ij} - \bar{y}^*_{ij}}{SD(y^*_{ij})},$$

where y_{ij} is the observed aid per person, $\bar{y}^*_{ij}$ is the mean of the adjusted posterior predictive distribution of aid per person, and $SD(y^*_{ij})$ is the standard deviation of the adjusted posterior predictive distribution of aid per person. By standardizing this way, we can compare differences between the predictive means and the observed values relative to the variation of the predictive distributions, so that a large difference between the observed and predictive means is either dampened or enhanced by a large or small standard deviation, respectively. In Figure 3, darker orange values indicate that more aid is going to these locations than our efficient scenario suggests, while dark purple values indicate that less aid is going to these locations than our efficient scenario suggests. Looking at Agriculture, our efficient scenario suggests that more aid should be sent to the central districts – locations with more area cultivated than other districts, as suggested by our efficient scenario. For total donations per person within each district, many of the districts that indicate more need for aid are on the western side of Malawi.

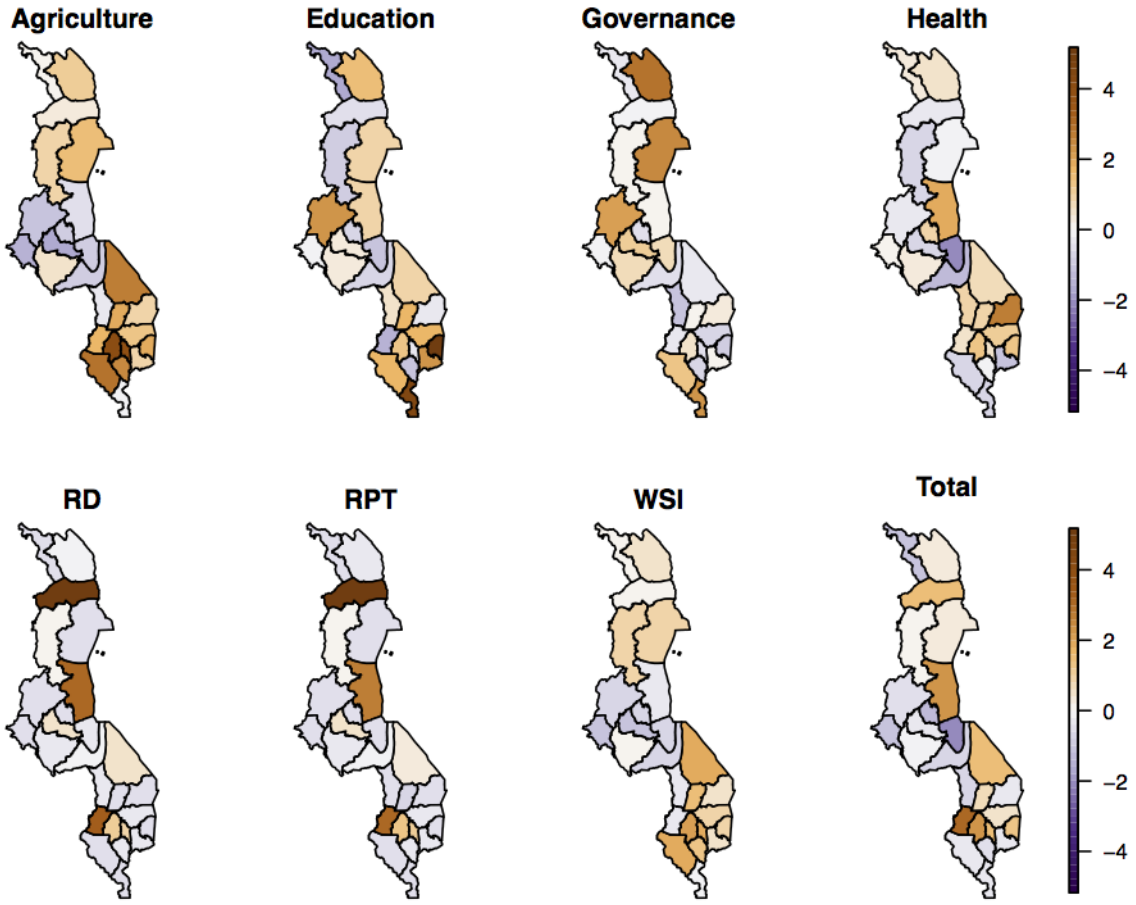

Figure 3: Standardized residual maps between the observed and efficient scenario donations per person for each sector and for all sectors combined (bottom right).

5 Conclusion

Foreign aid is important for many countries' economies and its allocation can be even more important to both the donors and citizens within the receiving country. But efficient allocation is a difficult problem and identifying whether or not its occurring is difficult too. Because it is so hard to determine allocation, there is a need for identifying ways to better allocate donations. In this paper, we made use of a unique dataset and examined how aid allocation within Malawi is related to various economic indicators within seven donation sectors. We found some evidence of efficient allocation, but on the whole, aid is not being allocated according to

economic indicators and could even be allocated in the opposite direction of what we might hope. We also used this model to develop a more efficient allocation of money according to a specific subset of the economic indicators and provided maps of adjustments to current aid allocation. Although most countries do not have databases with similar location-specific information on donations, these methods could be applied to examine efficiency of cross-national aid allocation.

Our model does not account for spatial confounding (see, e.g., Hughes and Haran, 2013). We found that the spatial random effect contributed in large part to the accuracy of the predictions in the models, and that the reduced-rank methods proposed for accounting for spatial confounding (Reich et al., 2006; Hughes and Haran, 2013; Hanks et al., 2015) could not capture that relationship as well, even when large numbers of bases were used. This is probably due to the small amount of data relative to the number of parameters required to fit such models, and requiring only one parameter for each sector to be estimated for the spatial dependence was ideal. Certainly, future methods could be explored to better account for spatial confounding.

Appendix

In this appendix we provide additional plots of the observed, predicted, and residual values of donations per person for each sector.

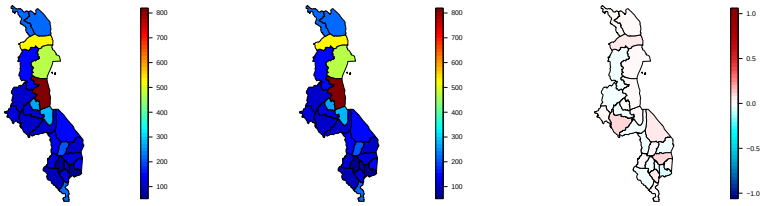

(a) Agriculture

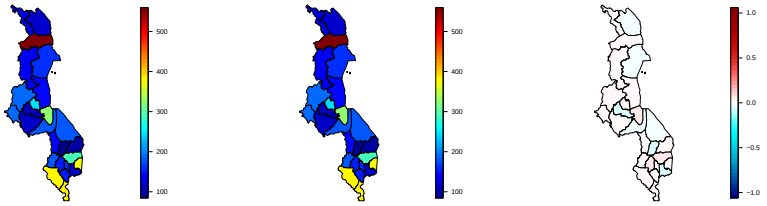

(b) Education

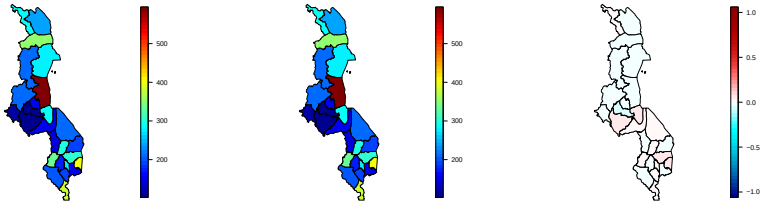

(c) Health

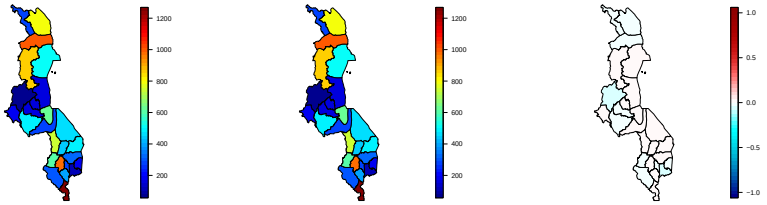

(d) Roads, Public Works and Transportation

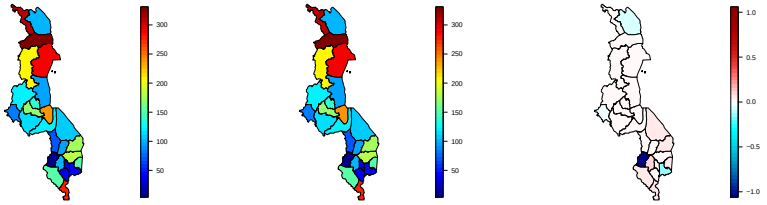

(e) Water, Sanitation, and Irrigation

Figure 4: Observed (left), predicted (center), and residual (right) donations plots for Agriculture, Education, Health, RPT, and WSI sectors. All values are in USD.

References

Aced, C. (2009). "Aid effectiveness in education." *Prospects*, 39, 107–108.

Alesina, A. and Dollar, D. (2000). "Who gives foreign aid to whom and why?" *Journal of Economic Growth*, 5, 33–63.

Banerjee, S., Gelfand, A. E., Finley, A. O., and Sang, H. (2008). "Gaussian predictive process models for large spatial data sets." *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 70, 4, 825–848.

Cornford, D., Csató, L., and Opper, M. (2005). "Sequential, Bayesian geostatistics: a principled method for large data sets." *Geographical Analysis*, 37, 2, 183–199.

De, R. and Becker, C. (2015). "The Foreign Aid Effectiveness Debate: Evidence from Malawi." *AidData Working Paper 6*. http://aiddata.org/aiddata-working-paper-series.

Diggle, P. J., Moyeed, R. A., and Tawn, J. A. (1998). "Model-based geostatistics (with discussion and rejoinder)." *Applied Statistics*, 47, 299–350.

Findley, M. G., Marineau, J. F., and Powell, J. (2016). "The swarm principle: A sub-national spatial analysis of aid targeting and donor coordination in sub-Saharan Africa." Unpublished Manuscript: University of Texas at Austin.

Findley, M. G., Powell, J., Strandow, D., and Tanner, J. (2011). "The Localized Geography of Foreign Aid: A New Dataset and Application to Violent Armed Conflict." *World Development*, 39, 11, 1995–2009.

Handcock, M. S. and Stein, M. L. (1993). "A Bayesian analysis of kriging." *Technometrics*, 35, 4, 403–410.

Hanks, E. M., Schliep, E. M., Hooten, M. B., and Hoeting, J. A. (2015). "Restricted spatial regression in practice: Geostatistical models, confounding, and robustness under model misspecification." *Environmetrics*, 26, 243–254.

Hughes, J. and Haran, M. (2013). "Dimension reduction and alleviation of confounding for spatial generalized linear mixed models." *Journal of the Royal Statistical Society, Series B*, 75, 139–159.

McCullagh, P. and Nelder, J. A. (1989). *Generalized Linear Models*. 2nd ed. CRC Press.

Minasny, B. and McBratney, A. B. (2005). "The Matérn function as a general model for soil variograms." *Geoderma*, 128, 3, 192–207.

National Statistical Office (2008). *2008 Population and Housing Census*. Zomba, Malawi: Government of Malawi. http://www.nsomalawi.mw/images/stories/data_on_line/demography/census_2008/Main%20Report/Census%20Main%20Report.pdf.

— (2015). *Malawi in Figures 2015*. Zomba, Malawi: Government of Malawi. http://www.nsomalawi.mw/images/stories/data_on_line/general/malawi_in_figures/Malawi%20in%20Figures%202015.pdf.

Nunnenkamp, P., Sotirova, A., and Thiele, R. (2015). "Do Aid Donors Specialize and Coordinate within Recipient Countries? The Case of Malawi." *AidData Working Paper 10*. http://aiddata.org/aiddata-working-paper-series.

OECD (2015). *OECD Factbook 2015: Economic, Environmental and Social Statistics*. OECD Publishing. DOI: 10.1787/18147364.

Peratsakis, C., Powell, J., Findley, M., Baker, J., and Weaver, C. (2012). *Geocoded Activity-Level Data from the Government of Malawi's Aid Management Platform*. Washington D.C.: AidData and the Robert S. Strauss Center for International Security and Law. http://aiddata.org/subnational-geospatial-research-datasets.

Reich, B. J., Hodges, J. S., and Zadnik, V. (2006). "Effects of residual smoothing on the posterior of the fixed effects in disease-mapping models." *Biometrics*, 62, 4, 1197–1206.

Weaver, C., Davenport, S., Baker, J., Findley, M., Peratsakis, C., and Powell, J. (2014). *Malawi's Open Aid Map*. The World Bank. http://wbi.worldbank.org/wbi/Data/wbi/wbicms/files/drupal-acquia/wbi/malawi_case_study2_0.pdf.